

Towards a canonical framework for designing agents to support healthcare organisations

John Fox¹, David Glasspool¹, Sanjay Modgil¹, Pancho Tolchinsky², Liz Black¹ and members of the ASPIC project.³

Abstract. We have explored the application of multi-agent systems to in a number of experimental applications in healthcare [1, 2, 3, 4]. Despite promising results, the design and implementation of the individual agents and their interactions in these studies varied from application to application, which has negative theoretical and practical consequences. The ASPIC project³ is aimed at understanding the nature and role of argumentation in agent systems and is providing an opportunity to explore a standard, or canonical⁴, agent architecture that has clear theoretical foundations and can be reused across applications. This initial report summarises the work to date and motivates and describes the canonical architecture. The properties of the ASPIC agent's primary components are summarised as a set of software signatures, which specify the pre-condition – post-condition (input-output) invariant of the component. These signatures provide general constraints on the agent model, while allowing components to be implemented with different behaviour to meet the demands of different application domains.

1 INTRODUCTION

Although the use of autonomous agents and multi-agent systems in healthcare applications is now quite well motivated, agent technologies themselves remain somewhat immature and from a theoretical point of view are often ad hoc. In this paper we report progress towards a standard, or canonical, agent model which is intended to be both theoretically well motivated and technically well defined. A canonical description of the model will facilitate component reuse and can serve as a standard against which alternatives can be compared.

Our starting point is the “domino” agent model [5]. The domino model can be viewed as an extension of the classical BDI agent which incorporates the ability to make decisions under uncertainty. The domino framework has been successfully applied in a variety of versions and on a range of healthcare applications.

The ASPIC project is providing an opportunity to strengthen the formal semantics of the domino model (especially the inference and decision-making components of the model) and to add additional capabilities that exploit techniques, notably dialogue and learning, that are needed for healthcare. The ASPIC project is concerned with developing foundations for a general theory of argumentation in the context of autonomous agents. The eventual aim of ASPIC is to develop a software platform which supports application developers in constructing agents that employ some or all of the argumentation components, together with more general cognitive functions that are necessary for agents to operate autonomously in complex, unpredictable environments like clinical service organisations. The latter include goal management, problem solving and planning.

In this paper we motivate the ASPIC agent model via a number of illustrative healthcare applications, and then show how it can be defined in a general and implementation-independent form which subsumes many agent models to be found in the literature. We finish by presenting a complete (but provisional) set of canonical signatures for the components of the agent model as a basis for discussion about the practicality and utility of such standards.

2 ILLUSTRATIVE HEALTHCARE APPLICATIONS

We begin in this section with a brief survey of four multi-agent healthcare systems which have, in different ways, informed our approach. In the next section we discuss some common features of these models, and go on to develop a generalised (“canonical”) agent with wider applicability.

Huang et al [1, 6] provide an early demonstration of a multi-agent system for a distributed application in which each agent was capable of reasoning and making decisions. The technical adequacy of this approach was demonstrated in an application for distributed breast cancer. Each agent was capable of making decisions based on an argumentation approach [7, 8], but possessed different expertise and therefore needed to seek assistance in certain circumstances. For example a “general practitioner” agent has broad but shallow knowledge of the patient while a “cancer specialist” agent has deeper but narrower expertise.

Black et al [9] have also demonstrated the technical feasibility of an agent-network in the breast cancer domain (figure 1) in which individual agents are able to follow simple guidelines and enact patient management plans as well as to take decisions (e.g. decisions to refer patients from a primary care unit to a genetics

1 Advanced Computation Laboratory, Cancer Research UK, London, England.

2 Universitat Politècnica de Catalunya, Spain.

3 www.argumentation.org

4 "The usual or standard state or manner of something"; for example, "the canonical way to organize a file system is as a hierarchy, with extensions to make it a directed graph" Wikipedia, April 2006.

clinic). Black has used this framework to demonstrate that certain decision-making and referral management strategies have substantial potential benefits in speeding up the journey of breast cancer patients through a number of clinical services. Black's PhD work is currently concerned with developing and formalising the performatives and dialog protocols required to support such a network.

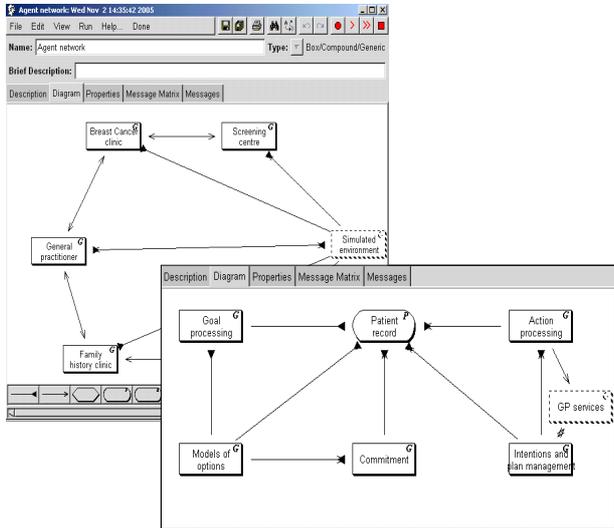


Figure 1. Black's [2] patient referral model showing the architecture of one agent (which is based on the "domino" agent model of Fox and Das [5]; see figure 4 below).

CARREL is being developed by Tolchinsky, Cortes and colleagues at the Technical University of Catalonia to support management of tissue and organ transplants. Although the National Transplant Service in Spain is among the world's best it is recognised that for a variety of reasons, many tissues and organs that are potentially available for transplantation to recipients are wasted ("discarded").

Among the reasons for this is that the decision to discard a donor's organ is currently based exclusively on the assessment of experts or 'agents' located at the site of the donor (the 'Transplant Coordinator' in figure 2). However, this ignores the fact that experts may disagree as to whether an organ is viable or not for transplantation. Hence, CARREL will support the identification of potential recipients of a donated tissue or organ across multiple transplant units, coordinate joint deliberation by donor and recipient agents through an argumentation-based dialogue, and evaluate the exchanged arguments. Through this procedure an organ that would ordinarily be discarded because it was considered non-viable by a donor agent may now be successfully transplanted if a recipient agent successfully argues that it is viable for the particular recipient that the agent represents.

The CREDO project is being developed at Cancer Research UK to support the "patient journey" for women with proven breast cancer and women at risk of contracting breast cancer. Figure 3 shows a schematic organisation diagram for a multidisciplinary care pathway for breast cancer. The diagram can be read from left to right as four phases of care: detection; work-up of diagnosis and treatment planning; enactment of treatment plan, and follow-up. The individual boxes represent collections of services provided by different clinical specialists, from the family doctor to the geneticist, medical oncologist, surgeon, radiologist, pathologist and palliative care professionals.

Although there is a rough left-to-right process of care members of

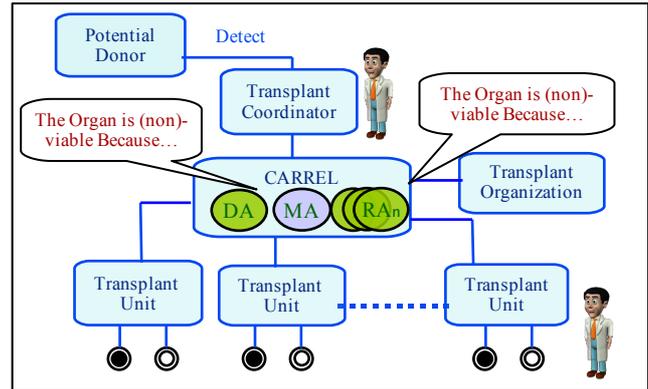


Figure 2. Illustration of the CARREL agent network in which a CARREL mediator agent (MA) directs and evaluates argument based deliberation between a donor agent (DA) representing the Transplant Coordinator and agents (RA) representing potential recipients in Transplant Units.

the clinical team do not provide services in a strict sequence but work in response to requests from colleagues, and in a continuing collaboration with those colleagues to achieve shared goals. The management of such care requires proper interpretation of events and effective communication; appropriate and safe decision-making; collaborative planning and timely execution of care plans, and the ability to avoid repeating mistakes.

The aim of the CREDO project is to create a multi-agent system that is capable of supporting the goals and processes of care at the individual and organisational levels (see video at <http://acl.icnet.uk/credo/CREDOwebsite/>). Each box represents an agent that is capable of providing specialist clinical services, applying capabilities for interpretation of situations, solving unexpected problems, making decisions about the best solutions and learning from experience.

A model of this domain suggests that there are around 65 significant decision points in the breast cancer pathway, all of which appear to be amenable to an argumentation based decision procedure [10]. In many cases these activities will require communication and sharing of responsibilities for clinical tasks with agents that have different specialist knowledge. These relationships are partially illustrated by the dataflow arrows shown in the model.

3 A CANONICAL AGENT MODEL

Although there are many ideas that are common to the different agent systems described above, notably the use of argumentation techniques, there are also significant variations.

Black's agent model incorporated abilities to represent and enact decisions and plans using the PROforma healthcare process specification language [5] while Huang's agent was limited to decision-making (based on the argumentation approach described in [7] and [8]). On the other hand both agent models supported a well-developed model of decision-making, involving raising goals on the basis of the agent's beliefs, generating possible solutions for those goals, a process of logical argumentation to determine the relative merits of alternative solutions to the same goal, and final decision on a particular solution (a new belief or a course of action). Both models also included a communication layer for communication between agents based on FIPA-like performatives. Neither of the communication models was able to access the logical argumentation services used during decision-making,

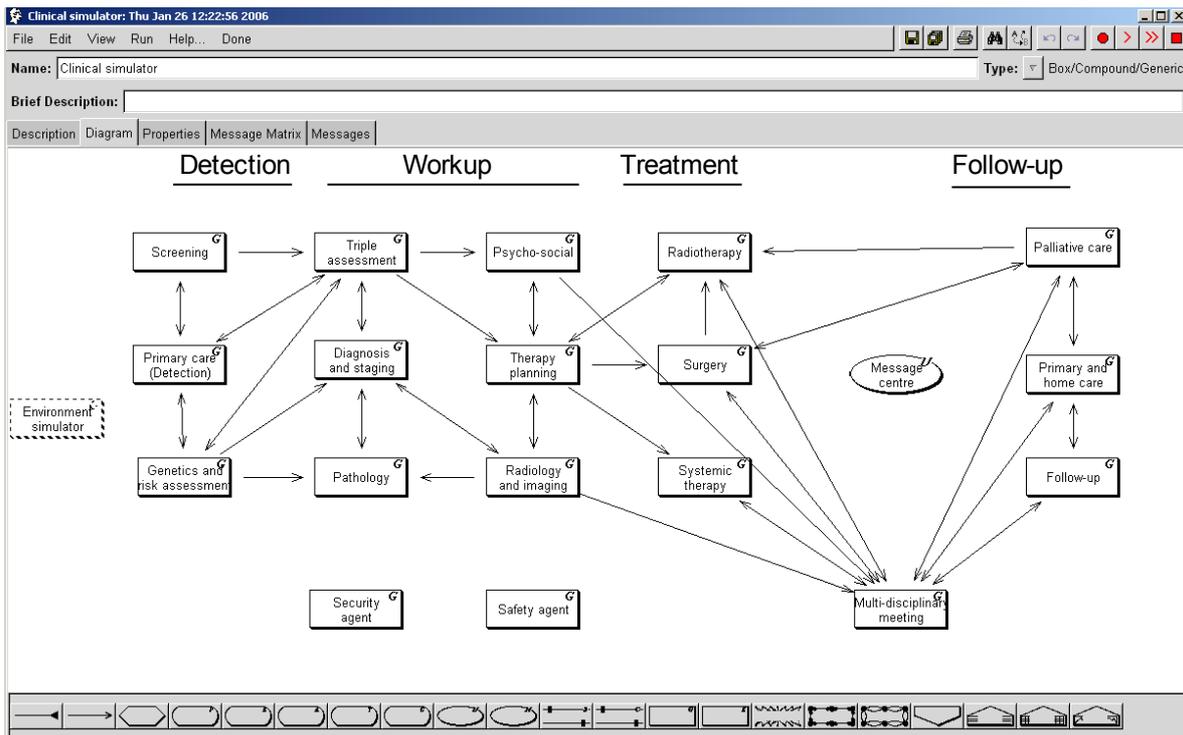


Figure 3. The CREDO organisation model for breast cancer; each box represents an agent which takes responsibility for supporting a particular collection of clinical and administrative services.

however, limiting the potential for collaborative activities. These problems are being addressed in the ASPIC project. Several ASPIC partners are developing and formalising inter-agent dialogue languages which incorporate performatives that facilitate coordination on collaborative tasks, such as joint decision-making or negotiation of services, where deliberative or dialectical argumentation *between agents* are required.

The model of an individual agent that is emerging from ASPIC brings together a number of different notions about agents that are established in the research literature, extended with argumentation roles (in inference, decision-making, learning and communication) and other capabilities. The aim is to develop an agent model that is based on a principled and well-defined set of component functions, and is generalizable to a wide range of real-world problems. The multi-agent systems discussed above have been used to inform and constrain the architecture, and the final two systems discussed, CREDO and CARREL, are also being used as test-beds for the emerging agent model.

In this section we summarise the main functions of this “canonical” agent. It is based on the “domino” agent architecture of Fox and Das [5] and we believe that it subsumes a range of proposals for integrated cognitive systems that have been described in a number of domains [11, 12].

The basic architecture can be viewed as an extension of the classic Belief-Desire-Intention (BDI) model of human practical reasoning developed by Bratman and others, and popular as a model of a rational software agent. Characteristically the BDI agent makes inferences over explicit mental states, including representations of its beliefs about present and predicted future states of the agent’s environment (e.g. [13, 14]. In our version of the BDI model we equate desires with goals, both epistemic goals (e.g. to find out

what is wrong with a patient) and practical goals (to decide what to do and do it) as discussed in [15]. Intentions are equated here with plans, which are viewed as commitments to carry out certain tasks to achieve goals. The term belief is used standardly to refer to what an agent holds to be true in the world (e.g. about a patient, the patient’s care, the other agents who are involved in that care) and what it believes about its own mental states (e.g. the justifications for its beliefs, goals and plans).

The domino model [5] (so-called because of its domino-like shape) extends the BDI framework to include problem-solving and decision-making functions (figure 4). Problem solving refers to processes by which an agent identifies possible solutions for goals (e.g. diseases that may be candidate diagnoses for a patient’s complaint or drugs that may be candidate treatments for the disease). In some cases problem solving may require the construction of complex models, as when an agent seeks an explanation of the aetiology of a familiar clinical presentation or a detailed plan of action to achieve a goal.

Decision-making may need to follow problem-solving if there is more than one candidate solution for a goal. In domains like medicine decisions may involve high levels of uncertainty about the true state of the environment or the consequences of actions on the environment. A canonical agent should be able to make the optimal choice under some definition of optimality. Traditionally, the optimal decision about a belief is the one that maximises the mathematical probability of being correct (about a diagnosis or level of clinical risk for example) and the optimal decision about action is the one that maximizes the expected utility of the effects of the action (e.g. benefits of treatment against costs and unwanted side-effects).

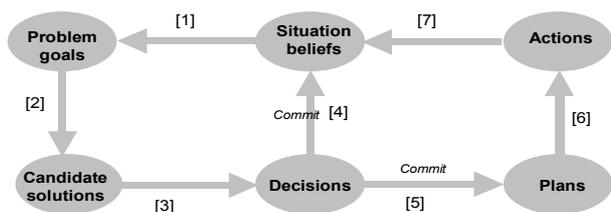


Figure 4. The “domino” framework for goal-based decision making and plan-based action [5]. Six types of information are acted on by seven processes: A set of beliefs can give rise to new goals (process 1) for an agent; a goal can be addressed (2) by determining candidate processes or plans for achieving it; candidates are assessed (3) and a decision is made on which new beliefs to hold (4) or plans to carry out (5); a plan is decomposed (6) into individual actions. Carrying out the actions may yield new beliefs (7).

It has often been observed that in healthcare the classical decision theoretical notion of optimality is frequently impractical (e.g. [5, 16]). A major factor here is the difficulty of determining usable probabilities and utilities in real-world situations. In areas like healthcare, where detailed statistical information on treatment outcomes for all possible combinations of patient characteristics is rarely available (or impossible to obtain in the case of new treatments) there are often qualitative or possibilistic arguments for or against an action but no quantitative probabilities or utilities. However, there is growing evidence that people do not naturally make decisions according to classical expected utility theory, although people are nonetheless surprisingly good at making decisions in difficult situations, under time pressure, or with incomplete information. ASPIC is therefore extending proposals for decision procedures based on argumentation, which are seen as more flexible and expressive (e.g. [16]). The domino agent includes a 3-step decision procedure from the point at which an agent raises a goal, viz: propose alternative solutions to the goal, formulate arguments for and against the alternatives, and assess the relative preferences over the options based on the strength and nature of the arguments. Theories of rational decision-making have a long history (e.g. [17, 18]) and creating something new or technically superior to conventional decision theory represents a significant challenge for agent technology:

“Ideally we would like to engineer agents that are as good at making decisions and acting upon them as we are. ... Unfortunately, like so many other problems encountered throughout the history of AI it has proven extraordinarily difficult to engineer agents that can select and execute good decisions for moderately complex environments. Nor is the problem likely to be solved in any meaningful sense for some time yet” (Wooldridge [14]).

We are more optimistic than Wooldridge on this matter having demonstrated that technologies based on the domino model work well on a range of non-trivial medical applications (reviewed in [19]). Two further extensions to the domino scheme are needed for applications being developed in the ASPIC project, inter-agent dialogue and learning.

The motivation for introducing communication and dialogue capabilities is clear on both theoretical and practical grounds. Wooldridge and Jennings [20] identified the ability to behave

socially as well as autonomously as key features of agent technology because in certain contexts an agent may only be able to achieve its goals by operating in concert with other agents.

The motivation for introducing learning capabilities into our agent model is particularly strong. From a practical point of view, for example, the UK National Health Service published “An organisation with a memory” [21] which noted that “adverse health care events cannot be eliminated from complex modern health care but [we can] ensure that lessons from the past are used to reduce the risk to patients in the future. The cost of adverse events is increasing; there is also a distressing similarity present in some of them. ... Specific types of adverse events are seen to repeat themselves at intervals, thus demonstrating that lessons have not been learned.”

The implications of this observation for the development of advanced clinical services is that wherever possible the services and/or service infrastructure should support organizational memory, maintaining records of what happened to patients, what decisions were taken and why, and what the outcomes for those patients were. This would not only support conventional clinical audit, in which clinicians critically review their own performance, but would also make it possible to apply machine learning techniques to update the system’s argumentation knowledge as it applies to clinical inference, decision-making, planning and so forth. In the context of the CARREL project case-based methods for learning from experience are being incorporated, and in the CREDO project rule induction methods developed by Mozina and Bratko in Ljubljana are being used, possibly in combination with case based methods as these represent complementary approaches.

3.1 Canonical components for agents

The extended agent architecture that is being developed within the ASPIC project is shown diagrammatically in figure 5, which shows the architecture within the COGENT simulation package [22] that is being used to prototype the ASPIC agent, and initial versions of the CARREL and CREDO applications. In this figure each rectangle is a component that implements one of the core functions outlined above, and each ellipse is a storage component which carries state information.

A number of increasingly complete and refined versions of this architecture have been implemented, testing the model’s adequacy on two medical scenarios. Provisional conclusions from this work are

- The agent model can be used in healthcare scenarios involving a small number of interacting agents (so far).
- The components appear to be sufficiently general to be effective across scenarios without major change to their capabilities. The primary caveats here are that considerably more experience is required to justify a firm statement along these lines.
- The components are modular with a uniform data interface between components. This facilitates component reusability and introduction of additional functions, such as dialogue and learning capabilities.

Our aim now is to consolidate the agent model in a way that captures the primary functions in a general and implementation-independent way. This “canonical component” model is intended to subsume many specific agent variants that have been proposed in the literature, providing clarity about the functions of the key components and a supportive framework for designing and implementing systems that fall within this general family of agent systems.

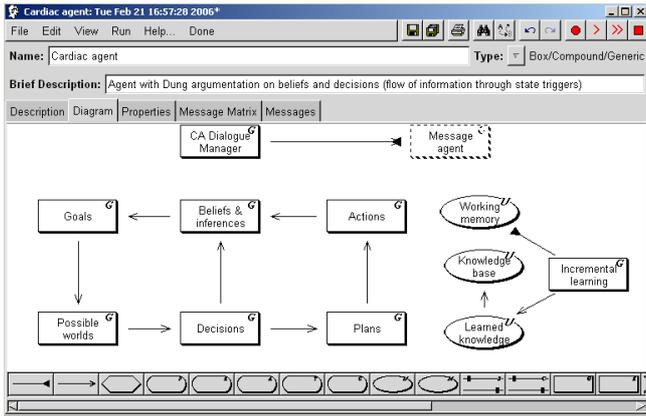


Figure 5. The ASPIC canonical agent mode implemented in COGENT. This can be seen as incorporating a classical BDI architecture is (beliefs and inferences; goals; plans) extended to include problem-solving (“possible worlds”) and decision-making, and a mechanism for implementing plans as actions on the world.

The domino is further extended with a dialogue manager for communication, and a learning component. Arrows represent the logical flow of information between components, though the actual flow is via working memory under the control of knowledge (COGENT memory components are shown as ellipses).

3.2 Component signatures for canonical agents

By taking an integrated view of the semantics of argumentation for four separate argumentation roles (inference, decision-making, dialogue and learning) ASPIC is making contributions to both argumentation theory and agent theory. As we have explained these roles are relevant to autonomous agents and multi-agent systems for healthcare applications and more generally. Despite this, however, there remain alternative points of view within the ASPIC consortium about the most appropriate or “correct” interpretations of these argumentation functions:

1. The ASPIC view of argumentation extends the traditional view of argumentation as defeasible inference [23] to ensure satisfaction of quality postulates. However it does not establish a relationship with classical approaches to inference under uncertainty (notably probabilistic inference) and in some cases may not be consistent with rational probability axioms. Healthcare applications depend strongly on notions like degree of belief and strength of evidence, and although we have argued above that in many cases possibilistic or qualitative, rather than quantitative probabilistic, inference is required, probabilistic reasoning is nonetheless important in many healthcare applications and this seems a significant area for further work.
2. The ASPIC framework for decision-making extends established argumentation approaches to decision making (e.g. [5]) but is restricted to decisions about action. In healthcare a substantial proportion of decisions are concerned with deciding what to believe as well as what to do.
3. Two distinct dialogue models have been developed in ASPIC. The first focuses on dialogues involving resolution of conflict, and provides a formal semantics for verifying the legality of moves in dialogues. The second model is more general in that it provides a semantics for dialogues that also involve deliberation. It

allows for less flexible dialogues than the first, but incorporates notions of strategy, such as choice of locution and locution content.

4. Finally, for learning, one group within ASPIC advocate case based reasoning to evaluate the relative strength of mutually attacking arguments, whereas another group use arguments to augment the classical machine learning process, and use classical machine learning to induce rules for argumentation. The two approaches appear complimentary and we expect that they might be used side by side.

At the very least there is more work to be done before finalizing a unique argumentation framework for use in agent systems. For complex domains like healthcare it may even be that a unique and sound formalization is impossible or impractical, because the demands of such domains entail radically different assumptions and tradeoffs, and possibly different logical axioms from situation to situation. In order to characterize a general ASPIC agent therefore, we are seeking to capture agent functions in a “canonical” form that abstracts from specific instantiations such as those identified in points 1-4 above.

To do this we have adopted the concept of a software signature from formal software engineering. A signature is defined by Spivey [24, p30] as “a collection of variable names, each with a type. Signatures are created by declarations, and they provide a vocabulary for making mathematical statements, which are expressed by predicates”. Table 1 presents a provisional proposal for a set of signatures for the core components of the domino agent model (1-8) augmented with signatures for dialogue and learning (9-12).

The signatures state general constraints on the preconditions and post-conditions of each component of the agent architecture, abstracting from details of the internal implementation⁵. In each case the typed variables above the line represent an input data pattern for the component, and the variable below the line is the type of the output of the component. At this level each component is a black box. It can be implemented in software or hardware; in the case of software it may be a conventional procedural algorithm, a set of production rules or a pure logic program. Following Spivey we assume that whatever the implementation technology the actual behaviour of the component can be described by an equivalent logic program.

The variables used in the signatures are as follows. **Belief** represents an agent’s beliefs about a situation, canonically viewed as propositions. **Goal** is a statement of an agent’s goals/desires, again represented propositionally, while **Option** is a candidate for achieving the goal, which may be a **Belief** or a **Plan**. **Arg** is an argument for or against an **Option**, and so on. Most of the signatures (and hence the software that instantiates them) refer to a **Theory** which is taken to be equivalent to a set of predicates in first-order logic, whatever the actual implementation. **Theory** subsumes a number of components (which we may distinguish in a

⁵ Wikipedia: “In computer programming, a precondition is a condition or predicate that must always be true just prior to the execution of some section of code or before an operation in a formal specification. Preconditions are sometimes tested using assertions within the code itself. Often, preconditions are simply included in the documentation of the affected section of code.”

1. Inference	2. Goals	3. Problem solving
Belief x Theory <hr/>	Belief x Theory OR Goal x Theory <hr/>	Goal x Belief x Theory <hr/>
Belief	Goal	Option
4. Deliberation	5. Argument aggregation	6. Decision taking: commitment
Goal x Option x Belief x Theory OR Goal x Option x Arg x Theory <hr/>	Goal x Option x Arg x Theory <hr/>	Goal x Option x Strength x Theory <hr/>
Arg	Strength	Belief OR Plan
7. Plan enactment	8. Action execution	9. Communication/dialogue
Goal x Plan x Theory <hr/>	Action x Resource <hr/>	Goal x Message x Theory <hr/>
Action OR Plan	Action OR Message	Tell OR Ask
10. Learning: MLBA	11. Learning: ABML	12. Learning: CBL
Belief x Goal x Theory <hr/>	Belief x Arg x Theory <hr/>	Belief x Option <hr/>
Theory	Theory	Theory

Table 1. Signature table for canonical components of an integrated agent based on the extended domino model.

final version of the model): the component mechanism (e.g. for inference or generating dialogue moves), its execution parameters (e.g. time or space restrictions on component operation) and domain ontology and other application-specific knowledge (e.g. knowledge required for a healthcare application).

There are several intended uses for such signatures. From a theoretical point of view each signature represents a claim that there is an invariant property or pattern that underlies a large (or at least interesting) class of agent designs. Thus, for example, signature 1 claims that deductive, defeasible and Bayesian inference are all constrained by the **Belief x Theory / Belief** signature, which is invariant: whatever process a particular implementation uses to implement this signature, it takes a belief (or set of beliefs) and a theory as input, and produces a new belief (or set of beliefs) as output. Problem-solving (signature 3) is defined as a process that determines a candidate belief or plan

given a set of beliefs, a goal and a theory. Deliberation (signature 4), the construction of arguments for decision-making, is a goal- and option-directed process (it takes a goal and an option, along with a set of beliefs and a theory, and constructs one or more arguments for or against the option), and aggregation (signature 5) takes a set of arguments for and against an option for addressing a goal, and determines an overall strength of support for the option. Communication acts (c.f. speech acts, dialogue moves) are simultaneously a specialization of general actions and strategic behaviours intended to achieve some communication goal: given a speech act and a communication resource a specific message to another agent may be formulated (signature 8), and given such a message, a communication goal and a dialogue theory, specific dialogue performatives may be enacted (signature 9). The three learning signatures (10, 11 and 12) apply to the three ways in which machine learning is being integrated with argumentation services within ASPIC: Machine-learning based argumentation

(MLBA) uses classical machine learning to induce rules for argumentation, and may be viewed as an augmentation to the deliberation process (signature 4); Argumentation-based machine learning (ABML) uses arguments to augment a classical machine learning process, and case-based learning techniques (CBL) uses case based reasoning to evaluate the relative strength of mutually attacking arguments, and thus augments the argument aggregation and commitment processes (signatures 5 and 6). The distinctive feature of the learning signatures, however, is that they specify that the output of the learning process is a (new or revised) theory, which may be used in turn within the other signatures.

From a more practical point of view signatures are classically seen as standard software-engineering formalisms which can guide and constrain software design (for example the design of agent software, or of specialized editing tools for capturing domain knowledge, e.g. goal editors, plan editors, decision editors). The other standard use for signatures is automatic verification of software design. If the component signatures are instantiated with a suitable logic program they can be subjected to formal verification techniques such as model checking.

4 CONCLUSIONS

Among ASPIC's goals are a formal semantics for argumentation functions for individual agents (including reasoning, decision-making and learning) and argumentation between agents in multi-agent institutions (including dialogues for shared decision making and negotiation). ASPIC is intended to provide a sound basis for developing practical algorithms for argumentation services in a variety of domains, including healthcare, and implemented components for reuse in practical applications. Applications to illustrate the practicality of argumentation technology are being developed for eBusiness (e.g. negotiation of credit), eGovernment (e.g. decisions about social benefits) and healthcare. The primary interest within healthcare is in multi-agent applications, including distributed decision making and negotiation in cancer care (the CREDO project) and transplant management (the CARREL project). The agent model described here is work in progress. It is intended to build links between ASPIC and existing work on autonomous agent architectures, exploring how argumentation services can be integrated into autonomous agents and multi-agent systems in a principled way.

The ASPIC project is providing an opportunity to extend and strengthen the domino agent model, both by adding capabilities, notably dialogue and learning, that are needed for healthcare applications, and by strengthening the formal semantics of the model, in particular in the areas of inference and decision-making where a clear and well-developed argumentation model is essential. Despite this progress it would be helpful to have a canonical description of the model to facilitate component reuse, and perhaps provide a standard agent model against which alternatives can be compared. In this paper we have discussed the motivation for these developments to the domino framework, and we have presented a provisional set of canonical signatures for the components of the ASPIC agent.

We expect the formalization of the agent model in terms of canonical signatures for its core computational functions to offer additional theoretical insights into the functions that are common to many agent theories and technologies, as well as providing a practical framework for designing and realizing complex multi-agent systems.

REFERENCES

- [1] J. Huang, N. Jennings and J. Fox. 'An agent based approach to health care management.' *Int. J. Applied Artificial Intelligence* 9 (4) (1995).
- [2] E. Black. 'Using Agent Technology to Model Complex Medical Organisations'. *AgentLink News* issue 12, April 2003.
- [3] P. Tolchinsk, S. Modgil and U. Cortes. Argument schemes and critical questions for heterogeneous agents to argue over the viability of a human organ. *Proc. AAAI Spring Symposium Series 2006* (AAAI Technical Report SS-06-01). (2006).
- [4] J. Fox, E. Black, D. W. Glasspool, S. Modgil, A. Oettinger, V. Patkar and M. Williams. 'Towards a general model for argumentation services'. *Proc. AAAI Spring Symposium Series 2006* (AAAI Technical Report SS-06-01). (2006).
- [5] J. Fox and S. Das. *Safe and Sound: Artificial Intelligence in Hazardous Applications*. AAAI and MIT Press. (2000).
- [6] J. Huang, N. Jennings and J. Fox. 'Co-operation in distributed medical care'. *Proc. 2nd Int'l Conference on Co-operative Information Systems*, 255-263, Toronto. (1994).
- [7] J. Fox, P. Krause and S. Ambler. 'Arguments, Contradictions and Practical Reasoning'. *Proc. ECAI 1992*, Vienna. (1992).
- [8] J. Huang, J. Fox, C. Gordon and A. Jackson-Smale. 'Symbolic decision support in medical care'. *Artificial Intelligence in medicine*, 5, 415-430 (1993).
- [9] E. Black and J. Fox. *Breast Cancer Referral System with Queuing Strategies*. Cancer Research UK, Advanced Computation Lab., Technical Report. (2002)
- [10] J. Fox, A. Alabassi, V. Patkar, T. Rose and E. Black. An ontological approach to modelling tasks and goals. *Comput Biol Med.* (2005).
- [11] D. W. Glasspool and R. Cooper. 'Executive Processes'. In R. Cooper (Ed.) *Modelling High Level Cognitive Processes*. New Jersey: Lawrence Erlbaum Associates. pp. 313-362. (2002).
- [12] D. W. Glasspool. 'The integration and control of behaviour: Insights from neuroscience and AI'. In D. N. Davis (Ed.) *Visions of Mind: Architectures for Cognition and Affect*. Hershey, PA: IDEA Group. pp. 208-234. (2005).
- [13] A. S. Rao and M. P. Georgeff. 'BDI Agents: from theory to practice'. *Proceedings of the first international conference on multi-agent systems (ICMAS-95)*, San Francisco. (1995).
- [14] M. Wooldridge, *Reasoning about rational agents*, Cambridge: MIT Press. (2000).
- [15] J. Fox, A. Alabassi, V. Patkar, T. Rose and E. Black. 'An ontological approach to modelling tasks and goals'. *Computers in biology and Medicine*, 2006 (in press).
- [16] D. W. Glasspool and J. Fox. 'Knowledge, argument and meta-cognition in routine decision-making'. In T. Betsch and S. Haberstroh (Eds) *The routines of decision making*. New Jersey: Lawrence Erlbaum. pp. 343-358. (2005).
- [17] D. V. Lindley. *Making Decisions* (2nd Edition). Wiley, Chichester, UK. (1985).
- [18] J. Fox. 'Logic, probability and the cognitive foundations of rational belief'. *Journal of Applied Logic*, 1, 1197-

2204. (2003).
- [19] J. Fox, V. Patkar, R. Thomson and E. Black. *PROforma guidelines and care pathways: evidence from trials and future plans*. This workshop.
- [20] M. Wooldridge and N. Jennings. 'Intelligent agents: theory and practice'. *The Knowledge Engineering Review*, 10(2), 115-152. (1995).
- [21] UK Department of Health. *An organisation with a memory*. Report of an expert group on learning from adverse events in the NHS, HMSO. (2000).
- [22] R. Cooper (Ed.) *Modelling High Level Cognitive Processes*. New Jersey: Lawrence Erlbaum Associates. pp. 313-362. (2002).
- [23] P. M. Dung. 'On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games'. *Artificial Intelligence* 77, 321-357. (1995).
- [24] J. M. Spivey. *The Z Notation: A Reference Manual*. Prentice-Hall, Hemel Hempstead, UK. (1989).